Customized Drive Electronics to Extend Silicon Optical Modulators to 4 Gb/s

Dean Samara-Rubio, *Member, IEEE*, Ulrich D. Keil, *Member, IEEE*, Ling Liao, Thorkild Franck, Ansheng Liu, Dexter W. Hodge, *Member, IEEE*, Doron Rubin, and Rami Cohen

***Abstract*—The data transmission bandwidth of a metal oxide semiconductor (MOS) capacitor Si optical modulator is extended from 1 to 4 Gb/s through the introduction of custom-designed low-impedance drive circuitry. Two distinct drive circuits were produced and tested—the first targeting 2.5 Gb/s data rate and 3 dB extinction ratio (ER), and the second having reduced voltage swing (1.3 V single-ended swing) while achieving an open eye at 4 Gb/s. The speed, power, and ER data collected are used to build a quantitative discussion of the challenges in achieving a power-efficient free-carrier modulator at bit rates above 1 Gb/s.**

***Index Terms*—BiCMOS integrated circuits, intensity modulation, optical modulation, optical waveguides.**

I. Introduction

THE BIT RATE of a metal oxide semiconductor (MOS) capacitor silicon optical modulator is extended from 1 to 4 Gb/s through the introduction of custom-designed low-impedance drive circuitry. This work follows upon Liu *et al.* [1] and Samara-Rubio *et al.* [2], which utilized the same silicon optical waveguide with embedded MOS capacitor (i.e., phase shifter). In [1], the phase shifter was shown to have small-signal electrical-to-optical conversion bandwidth of 2.5 GHz. In [2], a set of standard emitter-coupled logic (ECL) buffers was utilized to achieve 1 Gb/s bit rate with 5 dB extinction ratio (ER) and 2.9 W power dissipation. Based on trends of other modulator technologies, the 1 Gb/s data rate is significantly lower than might be expected for a device with 2.5-GHz small-signal conversion bandwidth and the power dissipation is higher than would be desirable based on specifications of driver circuits that are optimized for industry-standard modulators.

Certainly, the relative immaturity of high-speed silicon free-carrier modulators is at least partly the cause of these differences. A second factor is the relative inefficiency of the plasma dispersion effect compared to, for example, the quantum confined Stark effect utilized in electroabsorption modulators [3], [4]. However, the fundamental limit to the power required to charge and discharge the MOS phase shifter capacitors (presented in Section II-C) is more than an order of magnitude lower than the value achieved in [2], and, thus, a significant "gap" exists between the experimental data and the limiting case.

Though the preceding comparisons with existing modulator specifications can provide important insights, a broader view of the challenge to move data using light reveals a larger set of considerations. In particular, the emerging application of optics for ultrashort-distance very-high-bandwidth links (less than 100 m and hundreds of gigabits per second of aggregate bandwidth) demands: 1) cost reduction; 2) miniaturization; 3) photonic integration; and 4) embedded "intelligence." These four factors are driving many researchers to look for an appropriate manufacturing platform, of which silicon is one candidate.

Whether silicon is found to be a viable platform for photonic integration, and further, whether the free-carrier modulator will find a role is still to be determined. However, it must be emphasized that in existing applications as well as the emerging applications, minimizing power dissipation is often of paramount importance. Thus, understanding and significantly reducing the power dissipation "gap" is critical to understanding the performance, scalability, and applicability of the MOS free-carrier modulator.

In this paper, we focus on the Mach–Zehnder interferometer (MZI) design to explore the relationships between bit rate, power dissipation, and ER. Other promising free-carrier designs, such as those employing optical resonance [5] have been reported recently and can be analyzed in a similar manner as discussed here but have different characteristics and design tradeoffs and, thus, will likely require a different driver solution. Two custom-designed integrated circuits (ICs) have been developed and compared to the circuit in [2]. A simple metric, the bandwidth-to-power ratio (BWPR) is defined in (1a) to facilitate the comparison. The BWPR is interpreted at the power required per gigabit per second of data transmitted assuming 6 dB ER. The reference value of 6 dB is chosen as a reasonable ER for short-reach optical links that might employ an integrated silicon photonic chip. BWPR is primarily determined from eye diagrams collected using a digital communications analyzer (DCA) with a 20-GHz optical input. The DCA provides the ER and 20–80% rise and fall times (RT and FT). A nominal bit rate $(\mathrm{BR}^{\mathrm{nom}})$ is calculated from the RT and FT as shown in (1b). The power dissipation P is simply

Manuscript received February 3, 2005; revised July 14, 2005.

D. Samara-Rubio, L. Liao, A. Liu, and D. W. Hodge are with Intel Corporation, Santa Clara, CA 95054 USA (e-mail: dean.samara-rubio@intel.com; dean.sr@sbcglobal.net; ling.liao@intel.com; ansheng.liu@intel.com; hodge@ieee.org).

U. D. Keil and T. Franck are with Intel Copenhagen, Skovlunde DK-2740, Denmark (e-mail: ulrich.keil@intel.com; thorkild.franck@intel.com).

D. Rubin and R. Cohen are with Intel Corporation, Jerusalem 91031, Israel (e-mail: doron.rubin@intel.com; rami.cohen@intel.com).

Digital Object Identifier 10.1109/JLT.2005.859405

the product of supply voltage and supply current for the drive circuit plus modulator

$$\mathrm{BWPR} = \left(\frac{\mathrm{ER}}{6\ \mathrm{dB}}\right)\left(\frac{\mathrm{BR}^{\mathrm{nom}}}{P}\right) \tag{1a}$$

$$\mathrm{BR}^{\mathrm{nom}} = \left[3x\ \max\left\{\mathrm{RT}_{20}^{80}, \mathrm{FT}_{20}^{80}\right\}\right]^{-1}. \tag{1b}$$

Underlying the ER, and thus, the BWPR, is the efficiency of the free-carrier effect occurring in the optical waveguide. This relationship is discussed in Section II-A.

Both custom-designed drivers used in this work employ modified ECL output stages and SiGe heterojunction bipolar transistor (HBT). One of the drivers has been targeted to achieve large voltage swing of up to $2.5V_{\mathrm{pp}}$ at approximately 2.5 Gb/s (referred to as "Driver3G"), and the second has been targeted towards smaller voltage swing (up to $1.3V_{\mathrm{pp}}$) but with higher bandwidth (design targeting 8 Gb/s) and higher efficiency (referred to as "Driver8G"). To achieve a broadband low-impedance interface between the driver and modulator, they have been copackaged using a chip-on-board technique.

Section II below briefly reviews the fundamental operation of the active sections of the phase shifter. Small-signal impedance measurements of the phase shifter are reported and utilized to generate a compact equivalent circuit model that can be incorporated into circuit simulations. Section II concludes with a discussion of the fundamental limits to power dissipation and an outline of some reasonable expectations for power dissipation from a real drive circuit. Section III discloses design constraints, circuit topology, and experimental results for the Driver3G and modulator. Details of the circuit topology and internal structure are given. Section IV presents a similar treatment of Driver8G. To conclude, in Section V, the power efficiency of the drivers is tabulated, and a projection is made to determine the realistic performance targets for Driver8G combined with an improved modulator.

II. Phase Shifter Design and Characterization

The fundamentals of operation of the phase shifter and extensive experimental results have been presented in detail in [1] and [6]. In summary, the device is based on the free-carrier plasma dispersion effect, which can be exploited to create an electrically controllable optical path length for photons with energies below the bandgap energy. In silicon, the bandgap energy is approximately 1.1 eV, and this device is operated with photon energies of roughly 0.8 eV. In addition to altering the effective path length of the waveguide, the free-carriers absorb some of the light in the waveguide. This effect is manifested as a small voltage-dependent loss (VDL), which can be neglected for the majority of this discussion. The magnitude of VDL has been analyzed in [16] for a variety of MOS phase shifter waveguides.

A. Phase Shifter Efficiency

Fig. 1(a) shows a cross-sectional diagram of the phase shifter waveguide, and Fig. 1(b) shows the modeled transverse electric (TE) optical mode. Modeling and experiment have confirmed

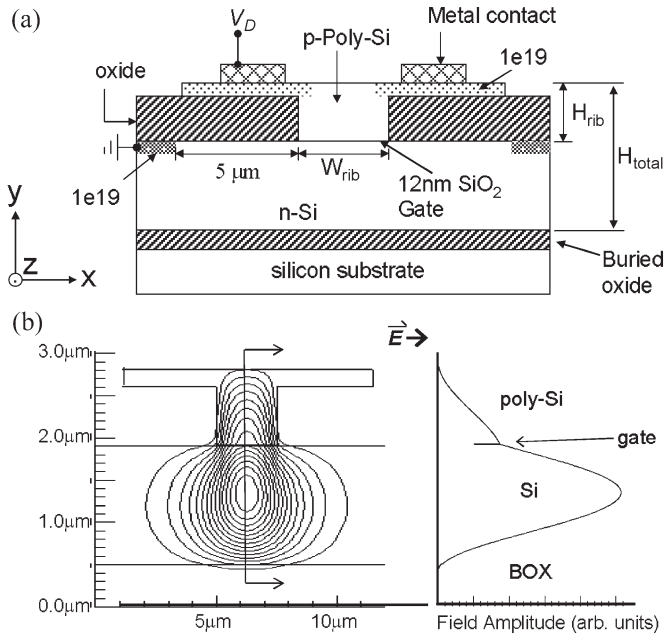

Fig. 1. (a) Cross-sectional view of the MOS waveguide phase shifter. (b) Modeled TE optical mode. Light propagates along the z-direction. Unless otherwise specified, measurements and discussions of phase shifters discussed in this paper have $H_{\mathrm{total}} = 2.3\ \mu$m, $\mathrm{W}_{\mathrm{rib}} = 2.5\ \mu$m, and $H_{\mathrm{rib}} = 0.9\ \mu$m.

that the mode shown in Fig. 1(b) is the single supported TE mode of the waveguide. Note that the discontinuity of the modal intensity near the gate is expected since the low-index gate dielectric (12 nm silicon dioxide) tends to repel the optical mode. Nevertheless, the waveguide described in Fig. 1 is capable of achieving a π phase shift in a length of 16 mm (i.e., L_π is 16 mm), assuming that the maximum electric field across the gate oxide is limited to approximately 5 MV/cm to ensure a suitably long gate dielectric lifetime. A precise value for maximum allowable field will need to be determined at a later stage of development [7].

Numerical integration of differential capacitance measurements [capacitance–voltage (CV) curves] gives the gate charge as a function of bias voltage, which can be compared to the measured phase shift. Fig. 2 shows differential capacitance (C) accumulation charge (Q) and phase (Φ) for an 8-mm-long phase-shift capacitor. C increases from left to right as the device transitions from depletion to accumulation. In the bias range of $+0.5$ to $+2$ V, C rises more gradually than expected from standard MOS models [8]. This is due to the presence of bias-dependent traps at or near the polysilicon–oxide interface combined with the fact that the surface potential of the polysilicon varies considerably more in these devices than in conventional MOS devices [9]. The polysilicon doping concentration near the gate is much lower than would be found in the conventional MOS devices [10].

The principal effect of this gradual transition is that the phase shift is expected to increase nonlinearly with voltage. In fact, Φ in Fig. 2 is slightly nonlinear with applied bias and can be fit to Q with proper selection of V_{fb} and Q_π. The best fit was obtained with $V_{\mathrm{fb}} = 0.9$ V and $Q_\pi = 500$ pF $\cdot$ V. Q_π is a relevant metric when designing the drive electronics whose goal is, in essence, the fast and efficient movement of charge.

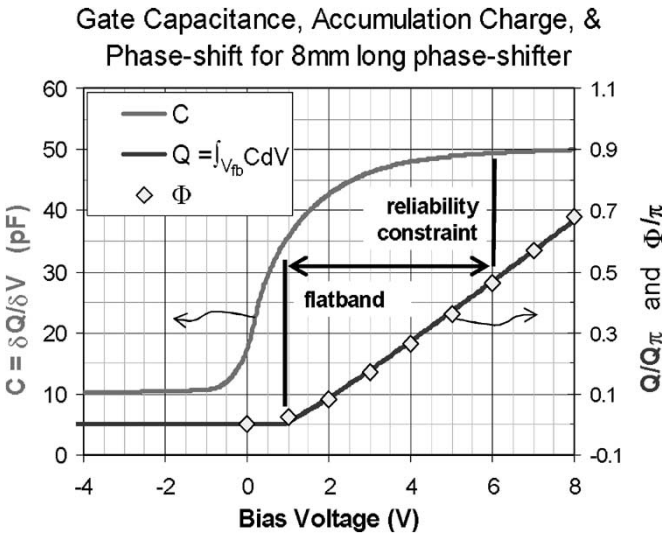

Fig. 2. Measured phase shift and gate capacitance for an 8-mm-long phase shifter with 12-nm gate oxide. The phase shift is in units of π radians. The second y-axis also includes the accumulation charge on the MOS gate, which is calculated as the numerical integration of the area under the CV curve starting at 0.9 V. A reliability constraint of 6.0 V is indicated as the maximum safe operating voltage for long lifetime operation.

Equations (2a)–(2d) shows that Q_π is already embedded in the BWPR metric and that for situations considered in this paper, BWPR is inversely proportional to Q_π. Equation (2a) expresses the output intensity of a push–pull-driven MZI as a function of the incremental phase shift $(\Delta\Phi)$ induced by the applied voltage and assumes that the device is biased around the quadrature point.

Equation (2a) also assumes that the optical splitter and combiner of the MZI are symmetric, and it excludes small corrections required to account for the voltage-dependent optical loss of the phase shifter. In (2b), the ER is expressed in terms of $\Delta\Phi$, and in the limit of small $\Delta\Phi$, the ER is shown to be proportional to $\Delta\Phi$. Equation (2c) is an expression of the proportionality of accumulation charge and phase, which is shown in Fig. 2 with a proportionality constant of π/Q_π. Finally, in (2d), (2b) and (2c) are combined with (1a) to show that the BWPR is inversely proportional to Q_π when ER is small. C, which appears in (2d), is the sum of the accumulation capacitances of the two arms of the MZI

$$\frac{I}{I_{\max}} = \frac{1+\sin(\Delta\Phi)}{2} \tag{2a}$$

$$\mathrm{ER} = 10\log\left[\frac{1+\sin(\Delta\Phi)}{1-\sin(\Delta\Phi)}\right] \approx 8.7(\Delta\Phi)|_{\Delta\Phi\leq\frac{\pi}{4}} \tag{2b}$$

$$\Delta\Phi = \left(\frac{\delta\Phi}{\delta Q}\right)\Delta Q = \left(\frac{\pi}{Q_\pi}\right)\frac{C_{\mathrm{gate}}V_{\mathrm{pp}}}{2} \tag{2c}$$

$$\mathrm{BWPR} \approx 0.73\pi\left(\frac{C_{\mathrm{gate}}V_{\mathrm{pp}}}{Q_\pi}\right)\left(\frac{\mathrm{BR}^{\mathrm{nom}}}{P}\right). \tag{2d}$$

The BWPR metric is most useful in situations in which the two expressions for ER given in (2b) agree to within 10% or better (i.e., for ERs of 6.5 dB or less).

Models and experiments have indicated that the value of Q_π is quite sensitive to the principal dimensions of the waveguide indicated in Fig. 1(a) as W_{rib} and H_{total}, but relatively less sensitive to the thickness of the gate oxide T_{ox}. (The value of H_{rib} is selected to ensure robust single-mode behavior.) For example, a 50% reduction in waveguide physical cross section, $W_{\mathrm{rib}} \times H_{\mathrm{total}}$ (from 2.5 μm $\times$ 2.3 μm to 1.65 μm $\times$ 1.8 μm) results in an approximately 60% reduction in Q_π, whereas a 50% reduction in T_{ox} from 12 to 6 nm has been found to give a small (15%) reduction in Q_π, because a thinner gate makes a smaller disturbance on the mode intensity distribution allowing more light to interact with the sheet charges. Clearly, reducing waveguide dimensions is a path to higher charge efficiency and allows the designer to move towards lower drive voltage, lower peak current requirement, higher ER, or smaller size.

However, as discussed in [16], one should expect an insertion loss penalty as waveguide dimensions are reduced. Thus, the net impact on optical signal-to-noise ratio (OSNR) as a result of scaling the waveguide dimensions must be evaluated for each specific implementation. Reduced OSNR will be manifested as higher bit error rate (BER) or lower maximum data rate for the link, or may be compensated with higher power consumption.

B. Bandwidth

In the previous section, it was indicated that the phase shift depends on the amount of charge on the gate. High-speed operation can only be achieved if this charge can be fully modulated in less than a bit period. It is the series resistance within the device, the output impedance of the driver, the gain bandwidth of the driver transistors, and the inductance of interconnect between them that are the potential limiters of bandwidth. In addition, even if the charge can be modulated sufficiently fast, the bandwidth may be limited if the electrical signal is not distributed to the phase shifter using a velocity-matching technique [11], [12]. The driver(s) reported here have been designed with a somewhat nonstandard velocity-matching technique that will be discussed in Section III.

For the purposes of understanding the fundamental operation of the device, short sections of the phase shifter were measured to determine the magnitude of the electrical-to-optical conversion (S21) and the frequency dependence of the input impedance (S11). The -3-dB bandwidth of S21 was reported in [1] for a 2.5-mm-long phase shifter to be 2.5 GHz. In this paper, we present S11 measurements of the phase shifter.

The S11 measurement is a purely electrical test that can be done on-wafer. It represents the electrical load placed upon the drive circuitry and, in fact, the primary motivation in measuring S11 is to provide an equivalent circuit model to guide the design of the output stage of the driver. A second motivation to measure S11 is to build a model that gives the quantity and distribution of charges on the gate capacitor as a function of frequency. From such a model, the S21 of a particular device might be predicted from a wafer-scale electrical measurement. This is an area of ongoing experimental and theoretical investigation.

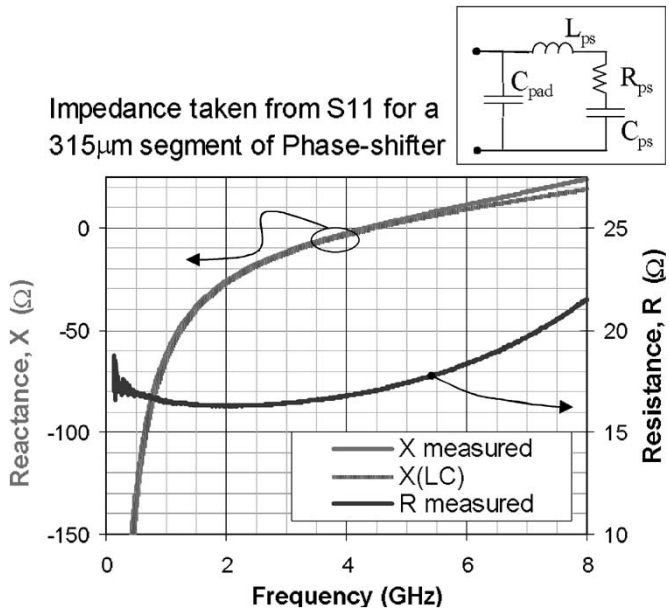

Fig. 3. Impedance taken from S11 measurement of a 315-μm-long segment of phase shifter. A compact model, shown in the inset, is used to fit the data.

Fig. 3 gives the resistance and reactance derived from the S11 of a 315-μm segment of phase shifter. The S11 measurement was performed using a ground–signal–ground (GSG) coaxial wafer probe and a vector network analyzer (VNA). To minimize transmission-line and distributed load effects, the contact pads are located at the midpoint of the device, and the device length is kept short. Also presented for comparison are the resistance and reactance of the compact model of the device, a simple inductance–capasitance–resistance (LCR) series circuit, which is shown in the inset of Fig. 3.

The contact pad capacitance $C_{\rm pad}$ is 0.11 pF. The values $C_{\rm ps} = 2.35$ pF and $L_{\rm ps} = 0.60$ nH are used to fit the measured data that include the gate capacitance plus all parasitic capacitances. $R_{\rm ps}$ is approximated to be $= 17\ \Omega$. The inductive component of the device is determined by the specific location of ground contacts and the routing of the metal traces on the die. In this case, the use of a single level of metal means that a C-shaped ground metal pattern surrounds the relatively long and narrow signal electrode $(315\ \mu{\rm m} \times 10\ \mu{\rm m})$. The resistance–capacitance (RC) cutoff frequency, $1/(2\pi RC)$ is in the range of 3.6 GHz, which is not in very good agreement with the S21 cutoff frequency of 2.5 GHz given in [1]. However, both pieces of data indicate that this phase shifter is capable of operation to 2.5 GHz or higher though previous attempts to use the phase shifter in a high-speed data modulator have achieved only 1 Gb/s operation. As discussed in [2], the primary reason for this discrepancy is that the large capacitances presented by long sections of phase shifter require low-impedance high-speed drive circuitry. Circuitry with appropriate specifications is not available in off-the-shelf components.

C. Lower Limit to Power Dissipation

The basic problem is to move charge on and off the capacitor plates, and the goal of this section is to use general arguments to set a reasonable lower limit on the power required to charge and discharge the modulator for data transmission. The electrostatic energy stored on the capacitor when transitioning from logic 0 to logic 1 state is given in (3) to be

$$U_0^1 = \frac{1}{2}C(V_1 - V_0)^2. \tag{3}$$

If the capacitor is charged by an ideal voltage step, then there will be an equal amount of energy dissipated as heat in the materials connecting the capacitor and source. If and when the voltage source is returned to the original voltage, the energy that was stored is dissipated. In this way, one cycle of charging and discharging the capacitor will dissipate CV^2 amount of energy.

In 1 s, the number of charge/discharge cycles is determined by the bit rate, data content, and encoding format. For a random binary data sequence, where a logic change is 50% likely for each bit transition, the number of charge/discharge cycles will average to BR/4.

In addition, note that the energy required to implement the ideal voltage source described above will be taken into account by considering additional gain stages in the driver. Any circuit technology will have a power gain associated with each stage and a value of $3X$ is a reasonable assumption for the stage gain (i.e., fanout). For example, assume that the voltage swing for each stage is the same, but the input capacitance reduces the factor of $3X$. The power for each driving stage can be calculated as above and added to the total and will give a factor of ~ 1.5 increase in power compared to (3), assuming more than two stages are needed. Finally, an additional factor of 2 is included to account for the fact that the push–pull modulator has two arms, and thus, two capacitors. Based on the above discussion, (4) provides the limiting case for power dissipation $P^{\rm min}$, written as

$$P^{\rm min} \cong 2 * (1.5)\left(\frac{\rm BR}{4}\right)C(V_1 - V_0)^2. \tag{4}$$

For the standard ECL circuitry of [2], the BR, C, and V are 1 Gb/s, 70 pF per arm, and $1.6V_{\rm pp}$ per arm to give a $P^{\rm min}$ of 134 mW. The experimentally measured value of 2.9 W is much larger than this limiting value. It is certainly not expected that a high-speed implementation would achieve the limiting value, but approaching 400% of the limiting value may be possible. To conclude this discussion, note that (4) indicates that an efficient driver would have power that increases linearly with bit rate, and this would seem to indicate scaling to 10 Gb/s or higher may be problematic. The question of whether the limiting value can approached, if not, why not, will be taken up in Section V.

D. Alternative Design: Transmission-Line Load

The challenge of quickly and efficiently modulating a capacitive load has been encountered in other types of optical modulators including lithium niobate and electro-absorption modulators (EAMs) [13], [14]. A common design approach in these devices is to form the capacitive load into a transmission

line that transforms the reactive (capacitive) behavior of the load into a resistive behavior. This transformation has two benefits. 1) It significantly reduces the frequency dependence of the power dissipation. 2) It provides the opportunity for velocity matching of the electrical and optical signal paths as was mentioned in Section I.

However, the capacitance per unit length of the MOS gate device tends to be much larger than is found in either of the aforementioned modulators, and this largely eliminates the benefits of a straightforward transmission-line approach. To see why this is true, consider the fundamental equations of transmission-line impedance Z_0 and wave propagation velocity $V_{\rm p}$

$$Z_0 = \sqrt{\frac{L}{C}} \tag{5a}$$

$$V_{\rm p} = \frac{2\pi}{\sqrt{LC}}. \tag{5b}$$

In (5a) and (5b), L and C are the inductance and capacitance per unit length of the transmission line. For the phase shifter shown in Fig. 1, these values are approximately 7.5 pF/mm and 1.0 nH/mm to give $Z_0 \sim 12\ \Omega$ and $V_{\rm p} \sim 7 \times 10^6$ m/s. For a Si core waveguide, the optical phase velocity of approximately 9×10^7 m/s is different from this calculated electrical phase velocity by a factor of 13. In fact, the large capacitive load given by the modulator inevitably leads to a solution with a prohibitively slow electrical wave, because, in practice, the inductance per unit length cannot be reduced without limit. The calculated impedance of 12 Ω is low, which leads to an equally low value for the terminating resistor(s) and generally higher power dissipation in comparison to a 50-Ω transmission line.

III. DRIVER3G DESIGN

A. Design Overview

The overall goal of the driver design was to provide robust circuit and packaging platform capable of testing present and future modulators with a compatible physical interface (as defined by number, size, and physical location of bond pads). The design is centered on a dual-drive MZI with 20–30 pF capacitance per arm for the radio frequency (RF) segments, a minimum peak-to-peak single-ended voltage swing of $2.0V_{\rm pp}$, and a target operating bandwidth of 2.5 Gb/s.

The physical interface between the driver and modulator was chosen to give low inductance while still allowing for a simple wire-bonding technique. As such, the RF phase shifters are 3.45 mm long and are divided into 11 equal sections of 315 μm. Each section has a bond pad for the two differential signals (applied to the p-type poly-Si gates of the phase shifters) and a bond pad for the RF return path (connected to the modulator n-type substrate). Fig. 4 is a photograph of a modulator and driver copackaged onto a printed circuit board (PCB).

The optical interferometer bias point is controlled by voltages taken directly from the PCB (connections seen on the lower side of the modulator in Fig. 4). This includes the modulator substrate direct current (dc) voltage, which is connected to the lowest potential in the system $V_{\rm EE}$.

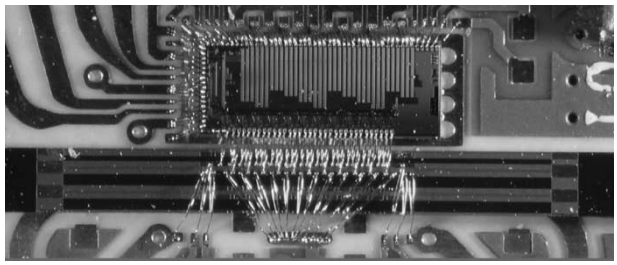

Fig. 4. Photograph of a driver die (top center) and modulator (bottom) copackaged on a PCB. The 11 buffers of the driver are stitched to 11 segments of the modulator phase shifters with short ($<$ 1 mm) wirebonds. Note that the modulator die was diced to be wider than needed and contains two MZIs, but only one is connected to the driver. The optical beam travels from left to right in this picture. The driver die measures 5×2 mm.

For the phase shifter of Fig. 1 (i.e., 2.5 μm $\times$ 2.3 μm waveguide with 12-nm gate oxide), $2V_{\rm pp}$ is expected to give approximately 2.5 dB ER. An alternate modulator, which would be compatible with this driver, is one based on a scaled waveguide with principal dimensions of 1.65 μm $\times$ 1.8 μm and gate oxide of 7 nm. The load capacitance is maintained in the target range, Q_π is reduced to 270 pF $\cdot$ V, and the expected ER is increased to $\sim$ 4.2 dB due to the higher charge efficiency.

Fig. 5 is a block diagram of the driver IC (left), the optical modulator (center), and optional electrical test load (right), which can be used in place of the MZI for electrical characterization. The driver IC is further subdivided into an input stage, denoted U1, a distribution stage, shown as an electrical transmission line, and an 11-segment output stage, denoted as U2_1 through U2_11. The optical modulator is represented by the set of 11 load ratio control (LRC) equivalent circuits found by S11 measurements (see Fig. 2).

A 70-GHz $f_{\rm T}$ SiGe HBT process was selected and current-mode logic (CML) was utilized for input stage of the driver IC. The array of 11 output cells is based on push–pull ECL architecture [15]. The circuit is powered from a single -5.2 V supply to meet or exceed the $2.0V_{\rm pp}$ minimum swing requirement across variations in supply voltage and temperature. The RF input to the driver is matched to a differential pair of 50-Ω PCB transmission lines with on-die termination resistors. DC control currents are used to set the output voltage swing and the bias current in the output cells.

To achieve optimal RT and FT, the data pulses are distributed to the 11 output cells via an internal transmission line. The TE optical mode of the waveguide has a group index calculated to be 3.69. The output voltage at segment 11 (i.e., the right side of Fig. 4) should be delayed by approximately 38 ps compared to the voltage at segment 1 (i.e., the left side of Fig. 4).

B. Characterization With Electrical Test Loads

The electrical test loads, shown on the right side of Fig. 5, are mounted in place of the optical modulator and wire bonded to the differential outputs of buffer stages U2_1 and U2_11.

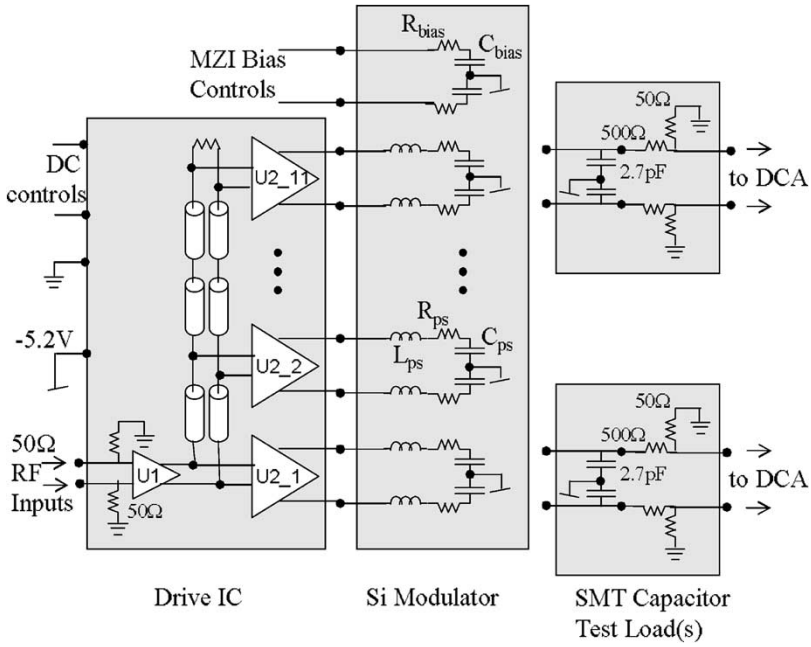

Fig. 5. Simplified circuit diagram of the 2.5-Gb/s SiGe HBT driver circuit shown adjacent to a matched push–pull MZI. On the right is an electrical test load that can be used in place of the MZI.

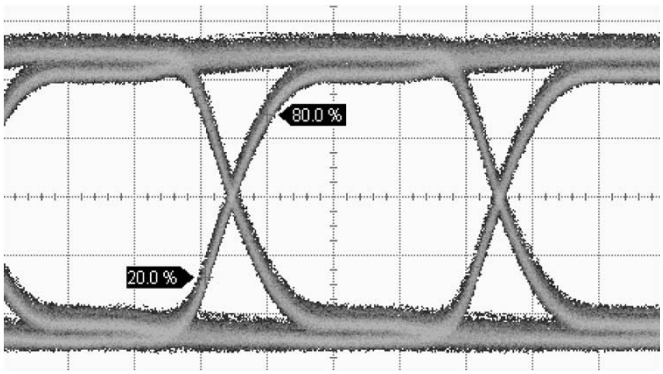

Fig. 6. Electrical eye diagram at 2.5 Gb/s for a single-ended buffer output with a 2.7-pF ceramic capacitor test load. The 20–80% RT is 93 ps. The RT and FT are approximately symmetric.

They allow for observation of three quantities that are difficult to access when the optical component is assembled: 1) propagation delay from buffer U2_1 to U2_11; 2) an indication for the highest bandwidth operation achievable from the driver when capacitively loaded; and 3) details of the single-ended voltage waveforms.

Each load is comprised of a pair of pancake-style surface mount ceramic capacitors and is positioned such that the wire bond from the driver to the capacitors is nominally the same length as the wire bond for the driver to phase shifters. The voltage on the test load is sampled through an additional wire bond to a 500-Ω surface mount resistor followed by a source-terminated 50-Ω transmission line, which is terminated at the electrical inputs of the DCA. A simple reactive test load was chosen over a resistor, because it best emulates the phase shifter.

The propagation delay between the outputs of buffers U2_1 and U2_11 is measured to be 26 ps. Following [11, eq. 3(b)], the 12-ps mismatch will give a 3-dB cutoff of nearly 40 GHz, which is much higher than the 3.6-GHz RC cutoff of the phase-shift capacitor. While the velocity match is not critical at this point in time, slight alteration in the internal transmission line layout or by a modest increase in the input capacitance of the buffer stages would improve the match.

Fig. 6 shows the measured eye diagram for a driver output into a 2.7-pF ceramic capacitor load, a power supply voltage of -5.2 V, and total average supply current of 0.73 A. The single-ended voltage swing is set at $2.1V_{\rm pp}$. Note that the measured 20–80% RT of 93 ps while the buffers are loaded means that this driver is capable of sourcing the necessary charge to the phase shifter to significantly exceed the 2.5 Gb/s target. From the RT, a peak switching current flowing into each of the capacitor loads is calculated to be 35 mA $(2.7\ \text{pF} * 2.1\ \text{V} * 0.6/93\ \text{ps})$.

A further inspection of Fig. 6 indicates that there is virtually no overshoot or ringing in the transitions and that the rising and falling edges are symmetric. As a result of this symmetry, the current sourced from the positive output of a differential buffer will equal the current drawn into the negative output of the same buffer, and thus, high frequency fluctuations in driver supply currents as well as the MZI substrate connection are an insignificant fraction of the switching current.

This fast, stable, and symmetric charging/discharging of a capacitive load are achieved through the use of the push–pull ECL architecture. A simplified schematic diagram of the output stage is presented in Fig. 7. The nodes VCS and VCAS are held at dc potentials for biasing the current source transistors Q3/Q3N and the cascode transistors Q2/Q2N, respectively. The differential high-speed data signal enters the circuit at nodes IN and INN, which are connected to the base terminals of transistors Q1/Q1N. The output voltage is taken from the emitter of Q1/Q1N, which tracks (follows) the voltage applied to the base but with a slightly lower potential due

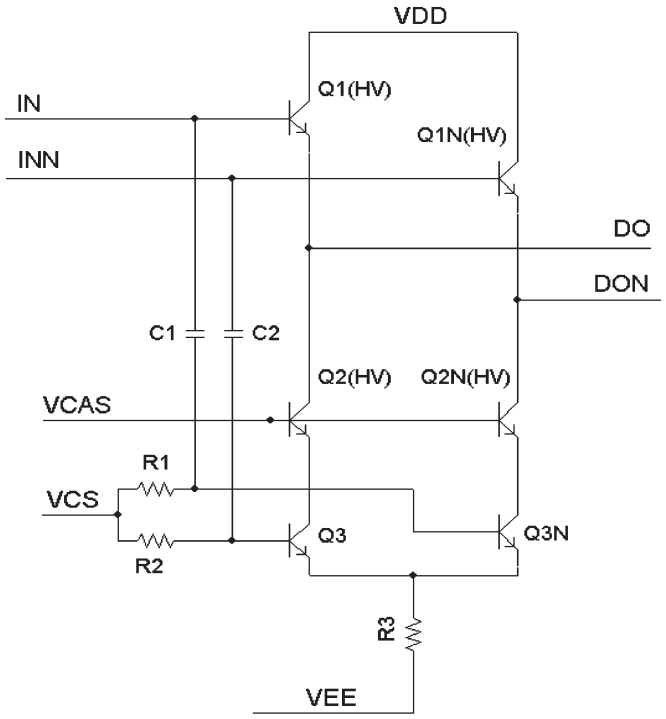

Fig. 7. Simplified circuit diagram of the output stage of the 2.5-Gb/s circuit Driver3G. The circuit is a differential push–pull emitter–follower in which the current source transistors (Q3/Q3N) are controlled through capacitive coupling to the inputs IN/INN. Nodes DO and DON are connected directly to the output bond pads.

to the base–emitter junction potential of the HBT. Transistors Q1/Q1N and Q2/Q2N have been annotated as "HV" in Fig. 7 to indicate that they are able to withstand higher $V_{\rm ce}$ compared to the standard HBT transistors provided by the process (e.g., Q3/Q3N). These HV transistors are needed to meet the $> 2V_{\rm pp}$ output swing requirement of the design and to ensure reliable performance across temperature and supply voltage variations. The penalty for the use of HV transistors is a reduction in switching speed.

The push–pull action occurs as the input signals are cross-coupled to the bases of the current source transistors Q3N/Q3 through the capacitors C1/C2. Optimal values of resistors R1/R2 and capacitors C1/C2 result in an approximate doubling of $I_{\rm Q3}$ during load discharge and reduction of $I_{\rm Q3}$ to near zero during load charging. In the case where the coupling capacitors are absent (or too small), the current sources pull constant current at the collector of transistors Q2/Q2N and will fight against the rising transition. During a falling transition, this current will limit the maximum discharge rate of the load capacitor. Thus, the correct device sizing minimizes contention and provides up to $4\times$ improvement in the available current to charge/discharge the load.

C. Optical Performance

As seen in the preceding section, Driver3G is capable of driving capacitive loads with RT/FT faster than required for 2.5 Gb/s data transmission (i.e., 93 ps achieved compared to the target value of 133 ps). In Section II, it was noted that the

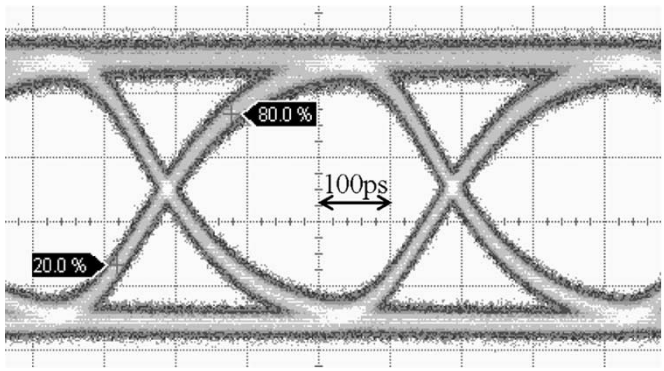

Fig. 8. Optical eye diagram using a 2.5-Gb/s non-return-to-zero (NRZ) pseudorandom bit sequence (PRBS) data stream into Driver3G copackaged with a modulator. The measured extinction ratio is 2.8 dB. The 20–80% RT is 152 ps, which is slower than the target value of 133 ps. Nevertheless, the eye is open and shows good SNR. The total power dissipation is 3.6 W.

resistance and inductance of the phase shifter and wirebonds will limit the flow of charge onto and off the capacitor and, thus, limit the bandwidth of the optical response. Fig. 8 contains the experimentally measured eye diagram for the optical response showing a 152-ps RT and an ER of 2.8 dB. The ER agrees reasonably well with the expected value of 2.6 dB ($2.1V_{\rm pp}$ applied to a pair of 3.45-mm phase shifters giving 0.19π of phase shift). In this measurement and all other optical measurements reported in this paper, the modulator is biased at the quadrature point by applying dc voltages to separate low-loss (lightly doped) waveguide sections in both arms of the interferometer.

The discrepancy between the 152-ps RT/FT of the optical signal compared to the 93-ps RT/FT of the electrical test load requires further investigation. To begin, note that neither the electrical eye diagram of Fig. 6 nor the optical eye diagram of Fig. 8 shows overshoot or ringing. Thus, despite the numerous active elements in the driver output stage, a single-pole damped exponential response is used to obtain an equivalent RC time constant from the 20%/80% RTs using

$$R_{\rm eq}C_{\rm load} \cong \frac{{\rm RT}_{20\%}^{80\%}}{\ln(0.8) - \ln(0.2)}. \tag{6}$$

For the electrical eye we obtain $R_{\rm eq}^{\rm electrical}$ of 25 Ω per buffer, and for the optical eye, we obtain an $R_{\rm eq}^{\rm optical}$ of 46 Ω per buffer (noting that each buffer is loaded with $C_{\rm ps} = 2.4$ pF from the S11 measurement). The 21-Ω difference between $R_{\rm eq}^{\rm electrical}$ and $R_{\rm eq}^{\rm optical}$ is primarily caused by the series resistance of the phase shifter $R_{\rm ps} = 17$ Ω, which was extracted from the S11 measurement. The wirebond length (and associated inductance) for the electrical and optical measurements are quite similar; thus, they have been taken into account.

The qualitative agreement between the S11 measurement and the eye measurements leads to the conclusion that the modulator resistance is the limiting factor for the optical transmission bit rate. The driver power dissipation with test loads is not significantly different than with the modulator, suggesting that there will be a significant improvement in BWPR by simply

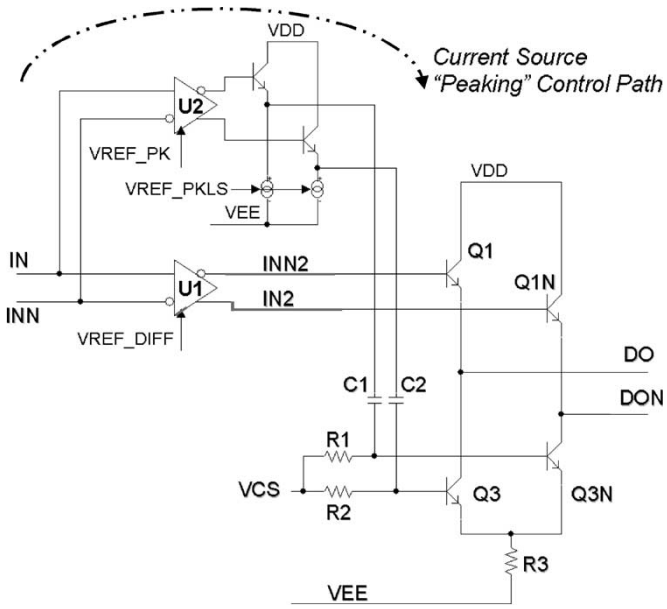

Fig. 9. For Driver8G, the basic push–pull ECL topology has been retained, but the cascode transistors have been removed, the transistors Q1/Q1N have been converter to standard HBTs, and limiting amplifiers U1 and U2 have been added to buffer and shape the input signal IN/INN.

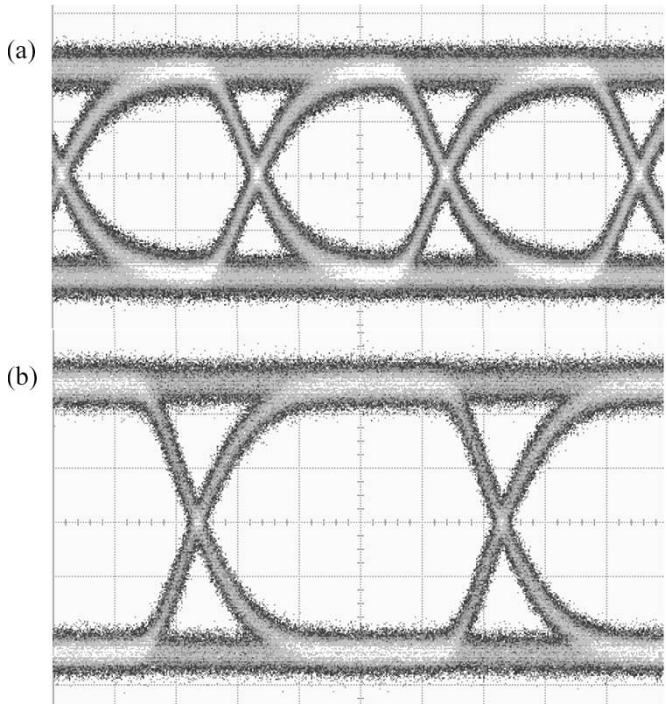

Fig. 10. Optical eye diagrams from Driver8G copackaged with a modulator. In (a), the input data rate is 4 Gb/s and the device has ER, RT, and P of 1.3 dB, 82 ps, and 2.7 W, respectively. In (b), the input data rate is 2.5 Gb/s, ER $=$ 1.7 dB, RT $=$ 91 ps, and only P $=$ 2.5 W.

replacing the modulator with one that has lower resistance and/or higher charge efficiency.

IV. Design for Higher Speed: Driver8G

The block diagram of Fig. 5 applies to Driver8G since only minor modifications were made to the input and distribution stages to accommodate the lower supply voltage of -3.3 V. Major modifications are to be found in the output stage. The large voltage swing and simplicity were sacrificed with the goal of achieving maximum bandwidth from the 70-GHz f_{T} SiGe HBT process. Fig. 9 is a schematic diagram of this output stage. All "HV" transistors have been converted to standard HBTs, the cascode transistors were removed from the emitter–follower current source and the preceding limiting amplifier was split into two, shown here as U1 and U2, which can be separately controlled.

The bias current of amplifier U1 is controlled via input VREF_DIFF to allow adjustment of the voltage swing applied to the modulator. The bias current of amplifier U2 is controlled via input VREF_PK; this adjustment ultimately affects the edge rate of the output stage with minimal impact on the voltage swing. A level shifter is also added to the "peaking" path and, via input VREF_PKLS, effectively adjusts the drive impedance of buffer U2. In total, three control signals have been added to the output stage, which, along with the VCS input, allow for independent control of output voltage swing, output impedance, and current consumption.

The experimentally measured optical eye diagrams from Driver8G are provided in Fig. 10(a) and (b). The horizontal and vertical scales are as follows: 80 ps/div and 60 μW/div for both eyes. The various control inputs were set to achieve the fastest stable edge rate while meeting the targeted values of 1.3 dB at 4 Gb/s in Fig. 10(a) and 1.7 dB at 2.5 Gb/s in Fig. 10(b). The RT of the driver is 82 ps at the lower ER and degrades slightly to 91 ps as the driver voltage swing is increased. However, the 10% degradation of edge rate was accompanied by a similar reduction in power supply current from 0.83 to 0.75 A. In total, the BWPR improves for Driver8G compared to Driver3G and is best at 1.7 dB ER and 91 ps edge rate.

Recently, we have tested improved optical modulator die based on improved manufacturing process to yield significantly lower R_{ps} for the same optical loss. The details of this work are available in [16]. The BWPR, ER, and other metrics are shown in the final row of Table I.

V. Conclusion on Power and Scalability

Table I is a summary of power, extinction ratio (ER), and bandwidth statistics of the optical data of [2] and eye diagrams found in Figs. 7 and 10. The final two columns of Table I are the calculated bandwidth-to-power ratio (BWPR) and the power efficiency defined as the actual power divided by the value given in (4). The table tells a mixed story. The introduction of the custom drivers has lead to a more than fourfold increase in measured optical bandwidth. However, achieving this bandwidth using the same generation of optical waveguide results in either an ER penalty or a power penalty. The net effect is that neither of the metrics, BWPR, or power efficiency has improved as much as the data rate.

It is expected that the power performance metrics can improve through reduction in modulator internal resistance R_{ps} and through the scaling down of waveguide dimensions. In fact,

TABLE I
PERFORMANCE SUMMARY OF DRIVERS. EXPERIMENTAL DATA ARE GIVEN IN ROWS 2–6. A COMPARISON OF ROWS 2 AND 4 INDICATES THAT DATA RATE HAS INCREASED FASTER THAN BWPR. ADDITIONAL MODULATOR SCALING (TO $\sim 1\ \mu$m $\times\ 1\ \mu$m WAVEGUIDES WITH ~ 4 nm T_{ox}) PLUS THE USE OF ADVANCED CMOS ARE POSSIBLE OPTIONS TO ACHIEVE THE 10-Gb/s TARGETS SHOWN IN ROW 1

Driver Design	20%-to-80% rise (fall) time (RT)	Nominal Bit Rate (BR^{nom})	Extinction Ratio (ER)	Power Dissipation (P)	Bandwidth-to-Power Ratio (BWPR)	Efficiency relative to P^{min} of Eqn. 2
	ps	Gbps	dB	W	Gbps/W @ 6dB ER	W/W
Target Values for 10Gbps	33	10.00	6.0	1	10.0	9%
Standard ECL [from Ref. 2]	415	0.80	5.0	2.7	0.2	4%
Driver3G	152	2.19	2.8	3.6	0.3	5%
Driver8G [@ 1.7dB ER]	91	3.66	1.7	2.5	0.4	6%
Driver8G [@ 1.3 dB ER]	82	4.07	1.3	2.7	0.3	3%
***Improved* Modulator** used with *Driver8G*	57	5.85	4.5	2.7	1.6	7%

aggressively scaled waveguides ($1\ \mu$m $\times\ 1\ \mu$m) can have $10\times$ lower Q_π than the current device. Such scaling combined with a thinner gate (~ 4 nm) may offer a path to meet the 10 Gb/s target values shown in row 1 of the table. Determining the appropriate driver technology and design approach will be vital to achieving the targets.

The SiGe heterojunction bipolar transistor (HBT) technology was a good choice for the first stage in development as it allows for flexible high-speed circuits with voltage swing > 1 V. However, with scaled optical devices, a voltage swing of 1 V and below may be acceptable and a digital complementary metal oxide semiconductor (CMOS) style of design may offer power efficiency advantages over HBT/emitter-coupled logic (ECL). For example, CMOS stages could be implemented with zero voltage headroom (i.e., they have rail-to-rail output swing) and the CMOS buffer would draw zero supply current during consecutive 1s and 0s. Furthermore, if sufficiently fast transistors are used to sustain the required bit rate and voltage swing, then the p-channel and n-channel transistors of the output stage could be driven in a "break-before-make" fashion, eliminating current flow directly from positive to negative power supplies. This type of implementation would be the most directly analogous to the hypothetical situation used to derive (4).

In summary, the data transmission bandwidth of a MOS capacitor Si optical modulator is extended from 1 to 4 Gb/s through the introduction of custom-designed low-impedance drive circuitry. In addition, a clear path is indicated to achieve 6 Gb/s or higher data transmission bandwidth through improvements in the optical modulator design. However, a power gap exists, which must be closed significantly to make the free-carrier modulator a compelling component in a silicon photonics platform.

ACKNOWLEDGMENT

The authors would like to thank B. Peterson, S. Koehl, S. Christensen, B. Danielsen, R. Lassen, M. Paniccia, and S. Pawlowski.

REFERENCES

[1] A. Liu, R. Jones, L. Liao, D. Samara-Rubio, D. Rubin, O. Cohen, R. Nicolaescu, and M. Paniccia, "A high speed silicon optical modulator based on a metal-oxide-semiconductor capacitor," *Nature*, vol. 427, no. 6975, pp. 615–618, Feb. 2004.

[2] D. Samara-Rubio, L. Liao, A. Liu, R. Jones, M. Paniccia, D. Rubin, and O. Cohen, "A gigahertz silicon-on-insulator Mach–Zehnder modulator," presented at the Optical Fiber Communication Conf., Los Angeles, CA, Feb. 22–24, 2004, PDP15.

[3] S. Schmitt-Rink, D. S. Chemla, W. H. Knox, and D. A. B. Miller, "How fast is excitonic electroabsorption?," *Opt. Lett.*, vol. 15, no. 1, pp. 60–62, Jan. 1990.

[4] V. Kaman, S. Z. Zhang, A. J. Keating, and J. E. Bowers, "High-speed operation of traveling-wave electroabsorption modulator," *Electron. Lett.*, vol. 35, no. 12, pp. 993–995, Jun. 1999.

[5] V. R. Almeida, C. A. Barrios, R. R. Panepucci, and M. Lipson, "All-optical control of light on a silicon chip," *Nature*, vol. 431, no. 7012, pp. 1081–1084, Oct. 2004.

[6] L. Liao, A. Liu, R. Jones, D. Rubin, D. Samara-Rubio, O. Cohen, M. Salib, and M. Paniccia, "Phase modulation efficiency and transmission loss of silicon optical phase shifters," *IEEE J. Quantum Electron.*, vol. 41, no. 2, pp. 250–257, Feb. 2005.

[7] H. Reisinger and R. Stengl, "Fundamental scaling laws of DRAM dielectrics," in *Proc. 3rd IEEE Int. Caracas Conf. Devices, Circuits and Systems*, Piscataway, NJ, 2000, pp. D26/1–D26/4.

[8] J. Hauser, *CVC NCSU Software, Version 5.0*, 2000, Raleigh, NC: Dept. Elect. Computer Eng., North Carolina State Univ.

[9] J. R. Hauser, private communication, Raleigh, NC: North Carolina State Univ.

[10] S. M. Sze, *Physics of Semiconductor Devices*, 2nd ed. New York: Wiley, 1981.

[11] R. Alferness, S. Korotky, and E. Marcatili, "Velocity-matching techniques for integrated optic traveling wave switch/modulators," *IEEE J. Quantum Electron.*, vol. QE-20, no. 3, pp. 301–309, Mar. 1984.

[12] R. C. Alferness, "Waveguide electrooptic modulators," *IEEE Trans. Microw. Theory Tech.*, vol. MTT-82, no. 8, pp. 1121–1137, Aug. 1982.
[13] G. L. Li, C. K. Sun, S. A. Pappert, W. X. Chen, and P. K. L. Yu, "Ultrahigh-speed traveling-wave electroabsorption modulator—Design and analysis," *IEEE Trans. Microw. Theory Tech.*, vol. 47, no. 7, pp. 1177–1183, Jul. 1999.
[14] E. L. Wooten, K. M. Kissa, A. Yi-Yan, E. J. Murphy, D. A. Lafaw, P. F. Hallemeier, D. Maack, D. V. Attanasio, D. J. Fritz, G. J. McBrien, and D. E. Bossi, "A review of lithium niobate modulators for fiber-optic communications systems," *IEEE J. Sel. Topics Quantum Electron.*, vol. 6, no. 1, pp. 69–82, Jan./Feb. 2000.
[15] V. G. Oklobdzija, "An ECL gate with improved speed and low power in a BiCMOS process," *IEEE J. Solid-State Circuits*, vol. 31, no. 1, pp. 77–83, Jan. 1996.
[16] L. Liao, D. Samara-Rubio, M. Morse, A. Liu, D. W. Hodge, D. Rubin, U. D. Keil, and T. Franck, "High speed silicon Mach–Zehnder modulator," *Opt. Express*, vol. 13, no. 8, pp. 3129–3135, Apr. 2005.

Dean Samara-Rubio (M'98) received the B.S. degree in electrical engineering from the University of Illinois at Urbana-Champaign, in 1991, and the Ph.D. degree in electrical engineering from the University of Texas at Austin, in 1997.

He is a Senior Researcher and Project Lead of the Corporate Technology Group, Intel Corporation, Santa Clara, CA.

Ulrich D. Keil (M'03–A'04–M'04) received the Ph.D. degree in physics from the University of Erlangen, Erlangen, Germany, and the postdoctoral degree from the Technical University of Denmark, Copenhagen, Denmark.

He is a Technology Officer in the Optical Networking-Components Division, Intel Copenhagen, Skovlunde, Denmark. Before joining Intel, he held research positions at Bell Labs, Murray Hill, NJ, at the Technical University of Eindhoven, Eindhoven, The Netherlands, and at the Technical University of Denmark.

Ling Liao received the B.S. and M.S. degrees in materials science and engineering from the Massachusetts Institute of Technology, Cambridge.

She is a Senior Researcher at the Corporate Technology Group, Intel Corporation, Santa Clara, CA.

Thorkild Franck received the Ph.D. degree in optical communications systems from the Danish Technical University, Kogens Lyngby, Denmark, in 1997.

He is a Design Validation Manager in the Optical Networking-Components Division, Intel Copenhagen, Skovlunde, Denmark. He characterizes phase-locked loops (PLL), equalizers, and broadband amplifiers for high-bandwidth optical systems.

Ansheng Liu received the Ph.D. degree from the University of Aalborg, Aalborg, Denmark, in 1992.

He is a Research Scientist Senior Staff developing silicon photonic devices in the Corporate Technology Group, Intel Corporation, Santa Clara, CA. Before joining Intel in 2000, he worked at NASA Ames Research Center, the National Institute of Standards and Technology, and at the University of Aalborg as an Assistant Professor. His interests include nonlinear optics of nanostructures, near-field optics, opto-electronics, and photonics. He has authored or co-authored more than 70 publications in these fields and contributed to 38 patents.

Dexter W. Hodge (M'05) received the B.S. degree in electrical engineering from the University of Rochester, Rochester, NY, and the M.S. degree in electrical engineering from the North Carolina State University, Raleigh.

He is a Graduate Technical Intern in the Photonics Technology Laboratory, Corporate Technology Group, Intel Corporation, Santa Clara, CA.

Doron Rubin received the B.S. degree in mechanical engineering and the M.B.A. degree from Ben Gurion University, Be'er Sheva, Israel.

He is a Process Engineer in the Photonics/MEMS Development Group, Technology Manufacturing Group, Intel Corporation, Jerusalem, Israel.

Rami Cohen received the B.S. and M.S. degrees in physical chemistry from Hebrew University, Jerusalem, Israel.

He is a Process Development Engineer in the Photonics Group, Technology Manufacturing Group, Intel Corporation, Jerusalem, Israel. Before joining Intel in 2004, he worked at Trellis-Photonics.